Search: The ACM Digital Library O The Guide

"data replication" "crash recovery" "atomic operation" "in-men



THE ACK DICITAL LIERARY

Feedback Report a problem Satisfaction survey

Terms used <u>data replication</u> <u>crash recovery</u> <u>atomic operation</u> <u>in</u> <u>memory index</u>

Found **81** of **201,890**

Sort results by Display

results

relevance expanded form

Save results to a Binder

Search Tips

Open results in a new

Try an <u>Advanced Search</u>
Try this search in <u>The ACM Guide</u>

window

Results 1 - 20 of 81

Result page: 1 2 3 4 5 next

Relevance scale

1 Research sessions: Research 20: Reliability: An integrated approach to recovery and high availability in an updatable, distributed data warehouse

Edmond Lau, Samuel Madden

September 2006 Proceedings of the 32nd international conference on Very large data bases - Volume 32 VLDB'2006

Publisher: VLDB Endowment

Full text available: pdf(669.84 KB) Additional Information: full citation, abstract, references, index terms

Any highly available data warehouse will use some form of data replication to tolerate machine failures. In this paper, we demonstrate that we can leverage this data redundancy to build an integrated approach to recovery and high availability. Our approach, called HARBOR, revives a crashed site by querying remote, online sites for missing updates and uses timestamps to determine which tuples need to be copied or updated. HARBOR does not require a stable log, recovers without quiescing the system ...

2 Disaster recovery techniques for database systems

•

Manhoi Choy, Hong Va Leong, Man Hon Wong November 2000 **Communications of the ACM**

Publisher: ACM Press

Full text available: pdf(412.04 KB) Additional Information: full citation, references, index terms

3 HFS: a performance-oriented flexible file system based on building-block



compositions

Orran Krieger, Michael Stumm

August 1997 ACM Transactions on Computer Systems (TOCS), Volume 15 Issue 3

Publisher: ACM Press

Full text available: pdf(383.87 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>, <u>review</u>

The Hurricane File System (HFS) is designed for (potentially large-scale) shared-memory multiprocessors. Its architecture is based on the principle that, in order to maximize performance for applications with diverse requirements, a file system must support a wide variety of file structures, file system policies, and I/O interfaces. Files in HFS are implemented using simple building blocks composed in potentially complex ways. This

Results (page 1): "data replication" "crash recovery" "atomic operat... Page 2 of 6 approach yields great flexibility, allowing an application ... **Keywords:** customization, data partitioning, data replication, flexibility, parallel computing, parallel file system Practical uses of synchronized clocks in distributed systems Barbara Liskov July 1991 Proceedings of the tenth annual ACM symposium on Principles of distributed computing PODC '91 Publisher: ACM Press Full text available: pdf(997.91 KB) Additional Information: full citation, references, citings, index terms Are quorums an alternative for data replication? Ricardo Jiménez-Peris, M. Patiño-Martínez, Gustavo Alonso, Bettina Kemme September 2003 ACM Transactions on Database Systems (TODS), Volume 28 Issue 3 **Publisher: ACM Press** Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> Full text available: pdf(2.23 MB) terms Data replication is playing an increasingly important role in the design of parallel information systems. In particular, the widespread use of cluster architectures often requires to replicate data for performance and availability reasons. However, maintaining the consistency of the different replicas is known to cause severe scalability problems. To address this limitation, quorums are often suggested as a way to reduce the overall overhead of replication. In this article, we analyze several qu ... **Keywords:** Data replication, availability, distributed transactions., quorums, scalability The cost of data replication Hector Garcia-Holina, Daniel Barbara October 1981 ACM SIGCOMM Computer Communication Review, Proceedings of the seventh symposium on Data communications SIGCOMM '81, Volume 11 Issue 4 **Publisher: ACM Press** Full text available: pdf(568.40 KB) Additional Information: full citation, abstract, references, index terms With the advent of data communication networks, researchers have been looking at the possibility of placing copies of a database at two or more nodes of a network. Such data replication is interesting because it makes the database accessible even when some of the nodes in the system fail. Furthermore, transactions which only read data may get faster access to the data when multiple copies exist. Due to the complex and time consuming update protocols, and due to the additional req ... MIDDLE-R: Consistent database replication at the middleware level Marta Patiño-Martinez, Ricardo Jiménez-Peris, Bettina Kemme, Gustavo Alonso November 2005 ACM Transactions on Computer Systems (TOCS), Volume 23 Issue 4 Publisher: ACM Press Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u>

Full text available: pdf(1.81 MB)

Additional Information: full citation, abstract, references, citings, index

The widespread use of clusters and Web farms has increased the importance of data replication. In this article, we show how to implement consistent and scalable data replication at the middleware level. We do this by combining transactional concurrency

Results (page 1): "data replication" "crash recovery" "atomic operat... Page 3 of 6 control with group communication primitives. The article presents different replication protocols, argues their correctness, describes their implementation as part of a generic middleware, Middle-R, and proves their feasibility with an extensive p ... **Keywords**: Database replication, eager data replication, middleware, scalability Replication in distributed systems: the Eden experience Jerre D. Noe, Andrew B. Proudfoot, Calton Pu November 1986 Proceedings of 1986 ACM Fall joint computer conference ACM '86 Publisher: IEEE Computer Society Press Full text available: pdf(1.36 MB) Additional Information: full citation, references, citings, index terms A Survey of Techniques for Synchronization and Recovery in Decentralized © Computer Systems Walter H. Kohler June 1981 ACM Computing Surveys (CSUR), Volume 13 Issue 2 Publisher: ACM Press Full text available: pdf(3.33 MB) Additional Information: full citation, references, citings, index terms 10 Persistent memory: a storage architecture for object-oriented database systems Satish M. Thatte September 1986 Proceedings on the 1986 international workshop on Object-oriented database systems OODS '86 Publisher: IEEE Computer Society Press Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> Full text available: pdf(1.13 MB) terms Object-oriented databases are needed to support database objects with a wide variety of types and structures. A persistent memory system provides a storage architecture for long-term, reliable retention of objects with rich types and structures in the virtual memory itself. It is based on a uniform memory abstraction, which eliminates the distinction between transient objects (data structures) and persistent objects (files and databases), and therefore, allows the same set of powerful and f ... 11 REQUEST II — a distributed database system for local area networks Marek Rusinkiewicz, Dimitrios Georgakopoulos November 1986 Proceedings of 1986 ACM Fall joint computer conference ACM '86 Publisher: IEEE Computer Society Press Full text available: pdf(1.02 MB) Additional Information: full citation, references, index terms 12 Conflict detection tradeoffs for replicated data Michael J. Carey, Miron Livny December 1991 ACM Transactions on Database Systems (TODS), Volume 16 Issue 4 Publisher: ACM Press Additional Information: full citation, references, citings, index terms, Full text available: pdf(2.50 MB)

<u>review</u>

Resu	ults (page 1): "data replication" "crash recovery" "atomic operat Page 4 o	f 6
	Keywords: concurrency control, replicated data	
13	Transaction management II: Nonblocking commit protocols Dale Skeen April 1981 Proceedings of the 1981 ACM SIGMOD international conference on	
	Management of data SIGMOD '81 Publisher: ACM Press Full text available: pdf(1.00 MB) Additional Information: full citation, abstract, references, citings	
	Protocols that allow operational sites to continue transaction processing even though site failures have occurred are called nonblocking. Many applications require nonblocking protocols. This paper investigates the properties of nonblocking protocols. Necessary and sufficient conditions for a protocol to be nonblocking are presented and from these conditions a method for designing them is derived. Both a central site nonblocking protocol and a decentralized nonblocking protocol are presented.	
14	Primitives for distributed computing Barbara Liskov December 1979 Proceedings of the seventh ACM symposium on Operating systems	
	December 1979 Proceedings of the seventh ACM symposium on Operating systems principles SOSP '79 Publisher: ACM Press	
	Full text available: pdf(877.17 KB) Additional Information: full citation, abstract, references, citings, index terms	
	Distributed programs that run on nodes of a network are now technologically feasible, and are well-suited to the needs of organizations. However, our knowledge about how to construct such programs is limited. This paper discusses primitives that support the construction of distributed programs. Attention is focussed on primitives in two major areas: modularity and communication. The issues underlying the selection of the primitives are discussed, especially the issue of providing robust beh	
15	Are wait-free algorithms fast?	
	Hagit Attiya, Nancy Lynch, Nir Shavit July 1994 Journal of the ACM (JACM), Volume 41 Issue 4	
	Publisher: ACM Press Full text available: 1 ndf(2.58 MR) Additional Information: full citation, abstract, references, citings, index	
	Full text available: pdf(2.58 MB) Additional information: <u>full citation</u> , <u>abstract, references</u> , <u>citings, index</u> <u>terms, review</u>	
	The time complexity of wait-free algorithms in "normal" executions, where no failures occur and processes operate at approximately the same speed, is considered. A lower bound of log n on the time complexity of any wait-free algorithm that achieves approximate agreement among n processes is proved. In contrast, there exists a non-wait-free algorithm that solves this problem in constant time. This implies an &OHgr(log n<	
	Keywords: approximate agreement, fault-tolerance, wait-free	
16	Scalable database replication through dynamic multiversioning Kaloian Manassiev, Cristiana Amza October 2005 Proceedings of the 2005 conference of the Centre for Advanced Studies on Collaborative research CASCON '05	

Full text available: pdf(313.41 KB) Additional Information: full citation, abstract, references, index terms

We scale the database back-end in dynamic content web servers on a set of database

Publisher: IBM Press

replicas while maintaining strong consistency. This is contrary to conventional wisdom in replicated databases which says that one could have either strong consistency or scalability, but not both. The key to scaling is a novel integrated fine-grained concurrency control and data replication algorithm called Dynamic Multiversioning that provides fine-grained distributed concurrency control at the level of a memory ...

	grantou distributou contain on a cita favor of a mismory m	
17 ②	WFS a simple shared file system for a distributed environment Daniel Swinehart, Gene McDaniel, David Boggs December 1979 Proceedings of the seventh ACM symposium on Operating systems	
	principles SOSP '79 Publisher: ACM Press	
	Full text available: pdf(751.34 KB) Additional Information: full citation, abstract, references, citings, index terms	
	WFS is a shared file server available to a large network community. WFS responds to a carefully limited repertoire of commands that client programs transmit over the network. The system does not utilize connections, but instead behaves like a remote disk and reacts to page-level requests. The design emphasizes reliance upon client programs to implement the traditional facilities (stream IO, a directory system, etc.) of a file system. The use of atomic commands and connectionless protocols n	
18	File servers for network-based distributed systems Liba Svobodova	
	December 1984 ACM Computing Surveys (CSUR), Volume 16 Issue 4	
	Publisher: ACM Press	
	Full text available: pdf(4.23 MB) Additional Information: full citation, references, citings, index terms, review	
19	Empirical performance evaluation of concurrency and coherency control protocols for	
	database sharing systems	
	Erhard Rahm June 1993 ACM Transactions on Database Systems (TODS), Volume 18 Issue 2	
	Publisher: ACM Press	
	Full text available: pdf(3.37 MB) Additional Information: full citation, abstract, references, citings, index terms, review	
	Database Sharing (DB-sharing) refers to a general approach for building a distributed high performance transaction system. The nodes of a DB-sharing system are locally coupled via a high-speed interconnect and share a common database at the disk level. This is also known as a "shared disk" approach. We compare database sharing with the database partitioning (shared nothing) approach and discuss the functional DBMS components that require new and coordinated solutions for DB-shar	
	Keywords : coherency control, concurrency control, database partitioning, database sharing, performance analysis, shared disk, shared nothing, trace-driven simulation	
20	Session 2B: Sorting and searching in the presence of memory faults (without	
	redundancy)	
•	Irene Finocchi, Giuseppe F. Italiano June 2004 Proceedings of the thirty-sixth annual ACM symposium on Theory of computing STOC '04 Publisher: ACM Proce	
	Publisher: ACM Press Full text available: Tondf(219.59 KB) Additional Information: full citation, abstract, references, index terms	

We investigate the design of algorithms resilient to memory faults, i. e., algorithms that, despite the corruption of some memory values during their execution, are able to produce a correct output on the set of uncorrupted values. In this framework, we consider two fundamental problems: sorting and searching. In particular, we prove that any $O(n\log n)$ comparison-based sorting algorithm can tolerate at most $O((n\log n)^{1/2})$ memory faults. Furthermore, we present one comparison-based sor ...

Keywords: combinatorial algorithms, memory faults, memory models, searching, sorting

Results 1 - 20 of 81

Result page: 1 2 3 4 5 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

<u>Terms of Usage Privacy Policy Code of Ethics Contact Us</u>

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



Search: The ACM Digital Library O The Guide

"data replication" "crash recovery" "atomic operation" "in-mem



THE ACK DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used data replication crash recovery atomic operation in memory index

window

Found **81** of **201,890**

Sort results by Display

results

Save results to a Binder

Search Tips

Open results in a new

Try an <u>Advanced Search</u>
Try this search in <u>The ACM Guide</u>

Results 21 - 40 of 81

Result page: previous 1 2 3 4 5 next

Relevance scale

Node autonomy in distributed systems

H. Garcia Molina, B. Kogan

January 2000 Proceedings of the first international symposium on Databases in parallel and distributed systems DPDS '88

Publisher: IEEE Computer Society Press

Full text available: pdf(925.19 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

The goal of this paper is to explore the notion of node autonomy in distributed computer systems. Some motivations for autonomy are exposed. Different facets of autonomy as well as relationships among them are discussed. Finally, we look into how autonomy affects other aspects of distributed computing, including timeliness, correctness, load sharing, data sharing, and data replication.

22 Evolving RPC for active storage

Muthian Sivathanu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

October 2002 ACM SIGPLAN Notices, ACM SIGOPS Operating Systems Review, ACM SIGARCH Computer Architecture News, Proceedings of the 10th international conference on Architectural support for programming languages and operating systems ASPLOS-X, Volume 37, 36, 30 Issue 10, 5, 5

Publisher: ACM Press

Full text available: pdf(1.56 MB) Additional Information: full citation, abstract, references

We introduce Scriptable RPC (SRPC), an RPC-based framework that enables distributed system services to take advantage of active components. Technology trends point to a world where each component in a system (whether disk, network interface, or memory) has substantial computational capabilities; however, traditional methods of building distributed services are not designed to take advantage of these new architectures, mandating wholesale change of the software base to exploit more powerful hardw ...

23 Soft updates: a solution to the metadata update problem in file systems

Gregory R. Ganger, Marshall Kirk McKusick, Craig A. N. Soules, Yale N. Patt May 2000 ACM Transactions on Computer Systems (TOCS), Volume 18 Issue 2

Publisher: ACM Press

Full text available: pdf(147.90 KB)

Additional Information: full citation, abstract, references, citings, index terms

Metadata updates, such as file creation and block allocation, have consistently been

identified as a source of performance, integrity, security, and availability problems for file systems. Soft updates is an implementation technique for low-cost sequencing of fine-grained updates to write-back cache blocks. Using soft updates to track and enforce metadata update dependencies, a file system can safely use delayed writes for almost all file operations. This article describes soft ...

24	Replicated distributed programs						
	Eric C. Cooper						
	December 1985 ACM SIGOPS Operating Systems Review, Proceedings of the tenth ACM symposium on Operating systems principles SOSP '85, Volume 19 Issue 5						
	Publisher: ACM Press						
	Full text available: pdf(1.12 MB) Additional Information: full citation, references, citings, index terms						
25	Occording the operational interest and according to the operational interest and according to the operational interest and according to the operation and a						
	database systems PODS '83						
	Publisher: ACM Press Full text available: pdf(1.28 MB) Additional Information: full citation, references, citings						
26	A distributed file service based on optimistic concurrency control						
②	Sape J. Mullender, Andrew S. Tanenbaum December 1985 ACM SIGOPS Operating Systems Review, Proceedings of the tenth ACM symposium on Operating systems principles SOSP '85, Volume 19 Issue 5						
	Publisher: ACM Press Full text available: pdf(910.04 KB) Additional Information: full citation, references, citings, index terms						
27	Transactional client-server cache consistency: alternatives and performance						
٩	Michael J. Franklin, Michael J. Carey, Miron Livny September 1997 ACM Transactions on Database Systems (TODS), Volume 22 Issue 3						
	Publisher: ACM Press						
	Full text available: pdf(452.41 KB) Additional Information: full citation, abstract, references, citings, index terms, review						
	Client-server database systems based on a data shipping model can exploit client memory resources by caching copies of data items across transaction boundaries. Caching reduces the need to obtain data from servers or other sites on the network. In order to ensure that such caching does not result in the violation of transaction semantics, a transactional cache consistency maintenance algorithm is required. Many such algorithms have been proposed in the literature and, as all provide the sam						
28 ③	MLR: a recovery method for multi-level systems David B. Lomet June 1992 ACM STCMOD Record - Proceedings of the 1993 ACM STCMOD international						
•	June 1992 ACM SIGMOD Record, Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92, Volume 21 Issue 2 Publisher: ACM Press						
	Full text available: pdf(1.10 MB) Additional Information: full citation, abstract, references, citings, index terms						
	To achieve high concurrency in a database system has meant building a system that						

copes well with important special cases. Recent work on multi-level systems suggest a systematic path to high concurrency. A multi-level system using locks permits restrictive low level locks of a subtransaction to be replaced with less restrictive high level locks when sub-transactions commit, enhancing concurrency. This is possible because sub-transactions can be undone via high level compensation actions ra ...

29 ②	Issues in the design and use of a distributed file system H. Sturgis, J. Mitchell, J. Israel	
	July 1980 ACM SIGOPS Operating Systems Review, Volume 14 Issue 3	
	Publisher: ACM Press Full text available: pdf(971.31 KB) Additional Information: full citation, abstract, references, citings	
	This paper discusses an independent file facility, one that is <i>not</i> embedded in an operating system. The <i>distributed file system</i> (DFS) is so named because it is implemented on a cooperating set of <i>server</i> computers connected by a communications, network, which together create the illusion of a single, logical system for the creation, deletion, and random accessing of data. Access to the DFS can only be accomplished over the network; a computer (or, more precisely, a program ru	
30	A survey of rollback-recovery protocols in message-passing systems E. N. (Mootaz) Elnozahy, Lorenzo Alvisi, Yi-Min Wang, David B. Johnson	
•	September 2002 ACM Computing Surveys (CSUR), Volume 34 Issue 3 Publisher: ACM Press	
	Full text available: pdf(549.68 KB) Additional Information: full citation, abstract, references, citings, index terms, review	
	This survey covers rollback-recovery techniques that do not require special language constructs. In the first part of the survey we classify rollback-recovery protocols into checkpoint-based and log-based. Checkpoint-based protocols rely solely on checkpointing for system state restoration. Checkpointing can be coordinated, uncoordinated, or communication-induced. Log-based protocols combine checkpointing with logging of nondeterministic events, encoded in tuples call	
	Keywords: message logging, rollback-recovery	
31 ③	Comparing data synchronization in Ada 9X and Orca Henri E. Bal January 1995 ACM SIGAda Ada Letters, Volume XV Issue 1	
	Publisher: ACM Press	
	Full text available: pdf(733.55 KB) Additional Information: full citation, abstract, citings, index terms	
	Protected object types are one of three major extensions to Ada 83 proposed by Ada 9X. This language feature is intended for light-weight data synchronization between tasks. The Orca parallel programming language has a very similar construct, the shared data-object, with which we have over five years of experience, both in usage and implementation. This paper compares protected objects and shared data-objects, with regard to design, usage, and implementation.	
32	Consistency in a partitioned network: a survey	
②	Susan B. Davidson, Hector Garcia-Molina, Dale Skeen September 1985 ACM Computing Surveys (CSUR), Volume 17 Issue 3	
	Publisher: ACM Press	
	Full text available: pdf(3.20 MB) Additional Information: full citation, abstract, references, citings, index	

terms, review

Recently, several strategies have been proposed for transaction processing in partitioned distributed database systems with replicated data. These strategies are surveyed in light of the competing goals of maintaining correctness and achieving high availability. Extensions and combinations are then discussed, and guidelines are presented for selecting strategies for particular applications.

33 ②	Read-only transactions in a distributed database Hector Garcia-Molina, Gio Wiederhold June 1982 ACM Transactions on Database Systems (TODS), Volume 7 Issue 2 Publisher: ACM Press	
	Full text available: pdf(2.16 MB) Additional Information: full citation, abstract, references, citings, index terms	
	A read-only transaction or query is a transaction which does not modify any data. Read-only transactions could be processed with general transaction processing algorithms, but in many cases it is more efficient to process read-only transactions with special algorithms which take advantage of the knowledge that the transaction only reads. This paper defines the various consistency and currency requirements that read-only transactions may have. The processing of the different classes of read	
	Keywords: R insularity, concurrency control, consistency, currency, query, read-only transaction, schedule, serializability, transaction, transaction processing algorithm	
34 ③	Language-level support for exploratory programming of distributed virtual environments Blair MacIntyre, Steven Feiner November 1996 Proceedings of the 9th annual ACM symposium on User interface software and technology UIST '96 Publisher: ACM Press	
	Full text available: pdf(1.68 MB) Additional Information: full citation, references, citings, index terms Keywords: distributed shared memory, distributed virtual environments, shared-data object model, virtual reality	
35 ③	Removing the overhead from software-based shared memory Zoran Radović, Erik Hagersten November 2001 Proceedings of the 2001 ACM/IEEE conference on Supercomputing (CDROM) Supercomputing '01 Publisher: ACM Press	
	Full text available: pdf(237.99 KB) Additional Information: full citation, abstract, references, citings, index terms	
	The implementation presented in this paperDSZOOM-WFis a sequentially consistent, fine-grained distributed software-based shared memory. It demonstrates a protocolhandling overhead below a microsecond for all the actions involved in a remote load operation, to be compared to the fastest implementation to date of around ten microseconds. The all-software protocol is implemented assuming some basic low-level primitives in the cluster interconnect and an operating system bypass functionality,	
36	Site optimal termination protocols for a distributed database under network	
③	partitioning David Cheung, Tiko Kameda August 1985 Proceedings of the fourth annual ACM symposium on Principles of	

Results (page 2): "data replication" "crash recovery" "atomic operat... Page 5 of 6 distributed computing PODC '85 Publisher: ACM Press Full text available: pdf(960.75 KB) Additional Information: full citation, references, citings, index terms Retrospection on a database system Michael Stonebraker June 1980 ACM Transactions on Database Systems (TODS), Volume 5 Issue 2 **Publisher: ACM Press** Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> Full text available: pdf(1.28 MB) terms This paper describes the implementation history of the INGRES database system. It focuses on mistakes that were made in progress rather than on eventual corrections. Some attention is also given to the role of structured design in a database system implementation and to the problem of supporting nontrivial users. Lastly, miscellaneous impressions of UNIX, the PDP-11, and data models are given. Keywords: concurrency, integrity, nonprocedural languages, protection, recovery, relational databases 38 JPernLite: an extensible transaction server for the World Wide Web Jack J. Yang, Gail E. Kaiser May 1998 Proceedings of the ninth ACM conference on Hypertext and hypermedia: links, objects, time and space---structure in hypermedia systems: links, objects, time and space---structure in hypermedia systems HYPERTEXT '98 Publisher: ACM Press Full text available: pdf(1.37 MB) Additional Information: full citation, references, citings, index terms Site optimal termination protocols for a distributed database under network partitioning David Cheung, Tiko Kameda July 1986 ACM SIGOPS Operating Systems Review, Volume 20 Issue 3 **Publisher: ACM Press** Full text available: pdf(826.50 KB) Additional Information: full citation, abstract, references, index terms Partition failure causes a major problem for the availability of a distributed database system. No protocol can consistently terminate all parts of a distributed transaction under all possible partitions, since some sites executing subtransactions may not know that other sites have already committed (or aborted) and, therefore, must wait until the failure is repaired. Under a site optimal termination protocol, the expected number of waiting sites is minimized, hence it maximizes the "availabilit ... A Self-Organizing Storage Cluster for Parallel Data-Intensive Applications Hong Tang, Aziz Gulbeden, Jingyu Zhou, William Strathearn, Tao Yang, Lingkun Chu

November 2004 Proceedings of the 2004 ACM/IEEE conference on Supercomputing SC '04

Publisher: IEEE Computer Society

Full text available: pdf(330.26 KB) Additional Information: full citation, abstract

Cluster-based storage systems are popular for data-intensive applications and it is desirable yet challenging to provide incremental expansion and high availability while achieving scalability and strong consistency. This paper presents the design and implementation of a self-organizing storage cluster called Sorrento, which targets data-intensive workload with highly parallel requests and low write-sharing patterns. Sorrento automatically adapts to storage node joins and departures, and the sys ...

Results 21 - 40 of 81

Result page: previous 1 2 3 4 5 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

<u>Terms of Usage Privacy Policy Code of Ethics Contact Us</u>

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



• The ACM Digital Library • The Guide Search:

"data replication" "crash recovery" "atomic operation" "in-mem



THE	ACM	DIRCHTAL	LIBRARY	**	۶.	7'	- 100	
. (3 6 4 2 432	127 1120 J. OCJ	Carried Color Color Color				6 W	

Feedback Report a problem Satisfaction survey

Terms used data replication crash recovery atomic operation in memory index

Found **81** of **201,890**

Sort results
by
Display

results

relevance expanded form

Save results to a Binder Search Tips

Try an Advanced Search Try this search in **The ACM Guide**

Open results in a new window

Results 41 - 60 of 81

Result page: previous 1 2 3 4 5 next

Relevance scale \square

The Farsite project: a retrospective



William J. Bolosky, John R. Douceur, Jon Howell April 2007

April 2007 ACM SIGOPS Operating Systems Review, Volume 41 Issue 2

Publisher: ACM Press

Full text available: pdf(225.09 KB) Additional Information: full citation, abstract, references, index terms

The Farsite file system is a storage service that runs on the desktop computers of a large organization and provides the semantics of a central NTFS file server. The motivation behind the Farsite project was to harness the unused storage and network resources of desktop computers to provide a service that is reliable, available, and secure despite the fact that it runs on machines that are unreliable, often unavailable, and of limited security. A main premise of the project has been that buil ...

Keywords: distributed debugging, formal system specification, project management, serverless distributed file system, software engineering, system design iteration, tech transfer

42 Concurrency and recovery for index trees

David Lomet, Betty Salzberg

August 1997 The VLDB Journal — The International Journal on Very Large Data

Bases, Volume 6 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(168.36 KB) Additional Information: full citation, abstract, citings, index terms

Although many suggestions have been made for concurrency in B \$^+\$-trees, few of these have considered recovery as well. We describe an approach which provides high concurrency while preserving well-formed trees across system crashes. Our approach works for a class of index trees that is a generalization of the B \$^{\rm link}\$-tree. This class includes some multi-attribute indexes and temporal indexes. Structural changes in an index tree are decomposed into a sequence of atomic actions, each one ...

Keywords: Access methods, B-trees, Concurrency, Indexing, Recovery

Session 2: Optimal termination protocols for network partitioning Francis Chin, K. V. S. Ramarao

Results (page 3): "data replication" "crash recovery" "atomic operat... Page 2 of 6

48

March 1983 Proceedings of the 2nd ACM SIGACT-SIGMOD symposium on Principles of database systems PODS '83

Publisher: ACM Press

Full text available: pdf(1.27 MB) Additional Information: full citation, abstract, references, citings

Commit protocols guarantee the consistency of distributed databases in absence of any failures. A commit protocol is resilient to a class of failures if it is possible to guarantee that a) databases at all operational sites in presence of these failures are consistent and b) other sites can be recovered consistently with these sites when the failure is repaired. Such a commit protocol is called *nonblocking* if no operational site needs to wait on a transaction which is incomplete at the ti ...

	transaction which is incomplete at the ti	
44 ③	A distributed persistent object store for scalable service Chao Jin, Weimin Zheng, Feng Zhou, Yinghui Wu October 2002 ACM SIGOPS Operating Systems Review, Volume 36 Issue 4 Publisher: ACM Press Full text available: pdf(1.03 MB) Additional Information: full citation, abstract, references	
	This paper presents a distributed persistent object store designed to simplify scalable service in cluster environment. This distributed object store, called TODS (Tsinghua Object Data Store), presents a single-imaged, transparent persistent and object-oriented view of the storage devices of the whole cluster. TODS is designed to be incremental scalable and efficient, and also has the properties of the high concurrency, high throughput and availability which are necessary for scalable service. T	
45 ②	T CHOITIANOC ANALYSIS OF SCACIAL DACK CHA database architectures	
	Full text available: pdf(1.54 MB) Additional Information: full citation, abstract, references, citings, index terms, review	
	The growing acceptance of database systems makes their performance increasingly more important. One way to gain performance is to off-load some of the functions of the database system to aback-end computer. The problem is what functions should be off-loaded to maximize the benefits of distributed processing. Our approach to this problem consisted of constructing several variants of an existing relational database system. INGRES, that partition the database system software into tw	
46 ②	Reliability Issues in Computing System Design B. Randell, P. Lee, P. C. Treleaven June 1978 ACM Computing Surveys (CSUR), Volume 10 Issue 2 Publisher: ACM Press	
	Full text available: pdf(3.95 MB) Additional Information: full citation, references, citings, index terms	
47	Checkpointing and rollback-recovery for distributed systems Richard Koo, Sam Toueg November 1986 Proceedings of 1986 ACM Fall joint computer conference ACM '86 Publisher: IEEE Computer Society Press	

Full text available: pdf(880.01 KB) Additional Information: full citation, references, index terms

Transactions and synchronization in a distributed operating system

Matthew J. Weinstein, Thomas W. Page, Brian K. Livezey, Gerald J. Popek December 1985 ACM SIGOPS Operating Systems Review, Proceedings of the tenth ACM symposium on Operating systems principles SOSP '85, Volume 19

Issue 5

Publisher: ACM Press

Full text available: pdf(974.32 KB) Additional Information: full citation, references, citings, index terms

49 Transactions and consistency in distributed database systems

Irving L. Traiger, Jim Gray, Cesare A. Galtieri, Bruce G. Lindsay

September 1982 ACM Transactions on Database Systems (TODS), Volume 7 Issue 3

Publisher: ACM Press

Full text available: pdf(1.48 MB)

Additional Information: full citation, abstract, references, citings, index terms

The concepts of transaction and of data consistency are defined for a distributed system. The cases of partitioned data, where fragments of a file are stored at multiple nodes, and replicated data, where a file is replicated at several nodes, are discussed. It is argued that the distribution and replication of data should be transparent to the programs which use the data. That is, the programming interface should provide location transparency, replica transparency, concurrency transparency, ...

Keywords: concurrency control, data partitioning, data replication, recovery

A workload characterization pipeline for models of parallel systems



William Alexander, Tom W. Keller, Ellen E. Boughter

May 1987 ACM SIGMETRICS Performance Evaluation Review, Proceedings of the 1987 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '87, Volume 15 Issue 1

Publisher: ACM Press

Full text available: pdf(961.78 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>

The same application implemented on different systems will necessarily present different workloads to the systems. Characterizations of workloads intended to represent the same application, but input to models of different systems, must also differ in analogous ways. We present a hierarchical method for characterizing a workload at increasing levels of detail such that every characterization at a lower level still accurately represents the workload at higher levels. < ...

Highly concurrent cache consistency for indices in client-server database systems



Markos Zaharioudakis, Michael J. Carey

June 1997 ACM SIGMOD Record, Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97, Volume 26 Issue 2

Publisher: ACM Press

Full text available: pdf(1.81 MB)

Additional Information: full citation, abstract, references, citings, index terms

In this paper, we present four approaches to providing highly concurrent B+-tree indices in the context of a data-shipping, client-server OODBMS architecture. The first performs all index operations at the server, while the other approaches support varying degrees of client caching and usage of index pages. We have implemented the four approaches, as well as the 2PL approach, in the context of the SHORE OODB system at Wisconsin, and we present experimen ...

52 HFS: a performance-oriented flexible file system based on building-block

	Compositions Orran Krieger, Michael Stumm May 1996 Proceedings of the fourth workshop on I/O in parallel and distributed systems: part of the federated computing research conference IOPADS '96 Publisher: ACM Press Full text available: pdf(1.83 MB) Additional Information: full citation, references, citings, index terms	
53	Adaptive algorithms for PASO systems Jeffery Westbrook, Lenore Zuck August 1994 Proceedings of the thirteenth annual ACM symposium on Principles of distributed computing PODC '94 Publisher: ACM Press Full text available: pdf(1.04 MB) Additional Information: full citation, references, index terms	
54	A coherent distributed file cache with directory write-behind Timothy Mann, Andrew Birrell, Andy Hisgen, Charles Jerian, Garret Swart May 1994 ACM Transactions on Computer Systems (TOCS), Volume 12 Issue 2 Publisher: ACM Press Full text available: pdf(3.21 MB) Additional Information: full citation, abstract, references, citings, index terms, review Extensive caching is a key feature of the Echo distributed file system. Echo client machines maintain coherent caches of file and directory data and properties, with write-behind (delayed write-back) of all cached information. Echo specifies ordering constraints on this write-behind, enabling applications to store and maintain consistent data structures in the file system even when crashes or network faults prevent some writes from being completed. In this paper we describe Keywords: coherence, file caching, write-behind	
55	Concurrency control in collaborative hypertext systems Uffe Kock Will, John J. Leggett December 1993 Proceedings of the fifth ACM conference on Hypertext HYPERTEXT '93 Publisher: ACM Press Full text available: pdf(1.05 MB) Additional Information: full citation, references, citings, index terms Keywords: collaborative work, concurrency control, distributed hypertext systems, events, extensibility, hyperbases, open architectures, supporting technologies, transaction	
56 ②	Michael Rabinovich, Edward D. Lazowska July 1992 Proceedings of the eleventh ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems PODS '92 Publisher: ACM Press Full text available: pdf(981.19 KB) Additional Information: full citation, abstract, references, index terms When failures occur during the execution of distributed commit protocols, the protocols may block in some partitions to avoid inconsistent termination of the transaction, thus	

Results (page 3): "data replication" "crash recovery" "atomic operat... Page 4 of 6

making data items in these partitions unavailable for accesses. We present a protocol that incorporates two new ideas with the goal of improving data availability. First, a new two-level voting scheme is proposed for deciding in which partitions to terminate the transaction. In this scheme, a choice is made based on t ...

57	Applications: Evaluation of database management systems for Erlang	
③	Emil Hellman September 2006 Proceedings of the 2006 ACM SIGPLAN workshop on Erlang ERLANG '06	
	Publisher: ACM Press Full text available: pdf(291.35 KB) Additional Information: full citation, abstract, references, index terms	
	Erlang/OTP's DBMS Mnesia is lacking in several important areas to consider when implementing very large databases with massive scalability requirements. This article reveals the result from a study examining what Erlang developers consider important aspects of DBMSs and an analytical hierarchy process (AHP) evaluation on four mature open source DBMSs based on those criteria. AHP is suggested as good method to evaluate DBMSs for Erlang projects. The criteria used in this evaluation were derived f	
	Keywords: AHP, DBMS, Erlang	
58 (Scalability issues in urban traffic systems David Al-Dabass, Evtim Peytchev, Mohamed Khalil, Manling Ren May 2006 Proceedings of the 1st international conference on Scalable information	
	systems InfoScale '06 Publisher: ACM Press Full text available: pdf(296.92 KB) Additional Information: full citation, abstract, references	
	Scalability in urban traffic systems comes in a wide variety of forms. At one end we have computing environments for building hierarchical traffic telematics distributed systems based on non-locking distributed shared memory algorithm. At the other end of scalability, to aid future automatic traffic movement in cities with displayed illuminated signs, algorithms were developed for the recognition of text, e.g. Chinese, character signs in real time which could ultimately form an important part of	
59 ③	Lazy replication: exploiting the semantics of distributed services Rivka Ladin, Barbara Liskov, Liuba Shrira	
	August 1990 Proceedings of the ninth annual ACM symposium on Principles of distributed computing PODC '90	
	Publisher: ACM Press Full text available: pdf(2.01 MB) Additional Information: full citation, references, citings, index terms	
60 ③	Research problems of decentralized systems with largely autonomous nodes Jerome H. Saltzer January 1978 ACM SIGOPS Operating Systems Review, Volume 12 Issue 1 Publisher: ACM Press Full text available: pdf(741.32 KB) Additional Information: full citation, references, citings	

Results 41 - 60 of 81 Result page: previous 1 2 3 4 5 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

Results (page 3): "data replication" "crash recovery" "atomic operat... Page 6 of 6

Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat Q QuickTime Windows Media Player Real Player



Search: The ACM Digital Library C The Guide

"data replication" "crash recovery" "atomic operation" "in-men

SEARCH

CT - LAYCTO	0.000	ADVICED A D	0.000000000	 	.38g.c	1	
INE	WCM	DIGITAL	LIBRARY		120		

Feedback Report a problem Satisfaction survey

Terms used <u>data replication</u> <u>crash recovery atomic operation in</u> <u>memory index</u>

window

Found **81** of **201,890**

Sort results by Display

results

relevance
expanded form

Save results to a Binder

Search Tips

Open results in a new

Try an <u>Advanced Search</u>
Try this search in <u>The ACM Guide</u>

Results 61 - 80 of 81

Result page: previous 1 2 3 4 5

Relevance scale

61 Advanced data processing in KRISYS: modeling concepts, implementation techniques, and client/server issues

Stefan Deßloch, Theo Härder, Nelson Mattos, Bernhard Mitschang, Joachim Thomas

May 1998 The VLDB Journal — The International Journal on Very Large Data Bases,

Volume 7 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(210.27 KB) Additional Information: full citation, abstract, citings, index terms

The increasing power of modern computers is steadily opening up new application domains for advanced data processing such as engineering and knowledge-based applications. To meet their requirements, concepts for advanced data management have been investigated during the last decade, especially in the field of object orientation. Over the last couple of years, the database group at the University of Kaiserslautern has been developing such an advanced database system, the KRISYS prototype. In this ...

Keywords: Client/server architectures, Consistency control, Object-oriented modeling concepts, Query processing, Run-time optimization

62 Scalable Networked Information Processing Environment (SNIPE)

Graham E Fagg, Keith Moore, Jack J Dongarra, Al Geist

November 1997 Proceedings of the 1997 ACM/IEEE conference on Supercomputing (CDROM) Supercomputing '97

Publisher: ACM Press

Full text available: pdf(77.42 KB) Additional Information: full citation, abstract, references, citings

SNIPE is a metacomputing system that aims to provide a reliable, secure, fault-tolerant environment for long-term distributed computing applications and data stores across the global InterNet. This system combines global naming and replication of both processing and data to support large scale information processing applications leading to better availablity and reliability than currently available with typical cluster computing and/or distributed computer environments.

Keywords: MetaComputing, RCDS, SNIPE, reliable, scalable, secure

63 <u>A replicated Unix file system</u>
Barbara Liskov, Robert Gruber, Paul Johnson, Liuba Shrira

Results (page 4): "data replication" "crash recovery" "atomic operat... Page 2 of 6 September 1990 Proceedings of the 4th workshop on ACM SIGOPS European workshop EW 4 Publisher: ACM Press Full text available: pdf(479.08 KB) Additional Information: full citation, references 64 Virtual time David R. Jefferson July 1985 ACM Transactions on Programming Languages and Systems (TOPLAS), Volume 7 Issue 3 Publisher: ACM Press Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.82 MB) terms, review Virtual time is a new paradigm for organizing and synchronizing distributed systems which can be applied to such problems as distributed discrete event simulation and distributed database concurrency control. Virtual time provides a flexible abstraction of real time in much the same way that virtual memory provides an abstraction of real memory. It is implemented using the Time Warp mechanism, a synchronization protocol distinguished by its reliance on lookahead-rollback, a ... 65 Consistency management in Deno Peter J. Keleher, Ugur Cetintemel December 2000 Mobile Networks and Applications, Volume 5 Issue 4 Publisher: Kluwer Academic Publishers Additional Information: full citation, abstract, references, citings, index Full text available: pdf(151.24 KB) terms We describe a new replicated‐ object protocol designed for use in mobile and weakly‐ connected environments. The protocol differs from previous protocols in combining epidemic information propagation with voting, and in using fixed per‐ object currencies for voting. The advantage of epidemic protocols is that data movement only requires pair‐ wise communication. Hence, there is no need for a majority quorum to be available and simultaneously connected at any single time. Th ... ⁶⁶ Fail-stop processors: an approach to designing fault-tolerant computing systems Richard D. Schlichting, Fred B. Schneider August 1983 ACM Transactions on Computer Systems (TOCS), Volume 1 Issue 3 **Publisher: ACM Press** Full text available: pdf(1.05 MB) Additional Information: <u>full citation</u>, <u>references</u>, <u>citings</u>, <u>index terms</u> Keywords: fail-stop Reliability mechanisms for SDD-1: a system for distributed databases Micael Hammer, David Shipman December 1980 ACM Transactions on Database Systems (TODS), Volume 5 Issue 4 Publisher: ACM Press Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> Full text available: pdf(3.06 MB) terms

This paper presents the reliability mechanisms of SDD-1, a prototype distributed database system being developed by the Computer Corporation of America. Reliability algorithms in SDD-1 center around the concept of the Reliable Network (RelNet). The RelNet is a

communications medium incorporating facilities for site status monitoring, event timestamping, multiply buffered message delivery, and the atomic control of distributed transactions. This paper is one of a series of compani ...

Keywords: atomicity, distributed databases, recovery, reliability

68	Fundamentals of fault-tolerant distributed computing in asynchronous environments	
	Felix C. Gärtner	
~	March 1999 ACM Computing Surveys (CSUR), Volume 31 Issue 1	
	Publisher: ACM Press Full text available: pdf(203.96 KB) Additional Information: full citation, abstract, references, citings, index terms, review	
	Fault tolerance in distributed computing is a wide area with a significant body of literature that is vastly diverse in methodology and terminology. This paper aims at structuring the area and thus guiding readers into this interesting field. We use a formal approach to define important terms like fault, fault tolerance, and redundancy. This leads to four distinct forms of fault tolerance and to two main phases in achieving them: detection	
	Keywords : agreement problem, asynchronous system, consensus problem, failure correction, failure detection, fault models, fault tolerance, liveness, message passing, possibility detection, predicate detection, redundancy, safety	
69	Problem space promotion and its evaluation as a technique for efficient parallel	
	computation	
•	Bradford L. Chamberlain, E. Christopher Lewis, Lawrence Snyder May 1999 Proceedings of the 13th international conference on Supercomputing ICS	
	'99 Publisher: ACM Press	
	Full text available: pdf(1.06 MB) Additional Information: full citation, references, citings, index terms	
70	Data-valued partitioning and virtual messages (extended abstract)	
③	Nandit Soparkar, Abraham Silberschatz April 1990 Proceedings of the ninth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems PODS '90	
	Publisher: ACM Press Full text available: pdf(1.48 MB) Additional Information: full citation, abstract, references, citings, index terms	
	Network Partition failures in traditional Distributed Databases cause severe problems for transaction processing. The only way to overcome the problems of "blocking" behavior for transaction processing in the event of such failures is, effectively, to execute them at single sites. A new approach to data representation and distribution is proposed and it is shown to be suitable for failure-prone environments. We propose techniques for transaction processing, concurrency control a	
71	Managing update conflicts in Bayou, a weakly connected replicated storage system	
(2)	D. B. Terry, M. M. Theimer, Karin Petersen, A. J. Demers, M. J. Spreitzer, C. H. Hauser December 1995 ACM SIGOPS Operating Systems Review, Proceedings of the fifteenth ACM symposium on Operating systems principles SOSP '95, Volume 29 Issue 5	
	Publisher: ACM Press	
	Full text available: pdf(1.56 MB) Additional Information: full citation, references, citings, index terms	

72 ②	A transputer T9000 family based architecture for parallel database machines Qiang Li, Naphtali Rishe December 1993 ACM SIGARCH Computer Architecture News, Volume 21 Issue 5 Publisher: ACM Press	
	Full text available: pdf(638.92 KB) Additional Information: full citation, abstract, index terms	
	Parallel computing is a promising way to achieve high performance in a database system. The disk access speed has been a well known bottleneck for database machines, and the data intensive nature and the random communication patterns of databases make the interconnection network in a database machine difficult to design. This article describes the design of a highly parallel, high throughput database machine based on the new T9000 transputer family and a large number of relatively inexpensive dis	
	Keywords : disk array, interconnection network, parallel database machine, semantic model, transputers	
73	Parallel algorithms column 1: models of computation Michael T. Goodrich	
	December 1993 ACM SIGACT News, Volume 24 Issue 4 Publisher: ACM Press	
	Full text available: pdf(434.71 KB) Additional Information: full citation, index terms	
74 ③	Database research at Wisconsin CORPORATE Univ. of Wisconsin March 1993 ACM SIGMOD Record, Volume 22 Issue 1 Publisher: ACM Press	
	Full text available: pdf(799.19 KB) Additional Information: full citation, index terms	
75 ③	Providing high availability using lazy replication Rivka Ladin, Barbara Liskov, Liuba Shrira, Sanjay Ghemawat November 1992 ACM Transactions on Computer Systems (TOCS), Volume 10 Issue 4 Publisher: ACM Press	
	Full text available: pdf(2.46 MB) Additional Information: full citation, abstract, references, citings, index terms, review	
	To provide high availability for services such as mail or bulletin boards, data must be replicated. One way to guarantee consistency of replicated data is to force service operations to occur in the same order at all sites, but this approach is expensive. For some applications a weaker causal operation order can preserve consistency while providing better performance. This paper describes a new way of implementing causal operations. Our technique also supports two other kinds of operations:	
	Keywords : client/server architecture, fault tolerance, group communication, high availability, operation ordering, replication, scalability, semantics of application	
76	A replicated Unix file system (extended abstract) Barbara Liskov, Robert Gruber, Paul Johnson, Liuba Shrira January 1991 ACM SIGOPS Operating Systems Review, Volume 25 Issue 1	

Results (page 4): "data replication" "crash recovery" "atomic operat... Page 4 of 6

Resu	lts (page 4): "data replication" "crash recovery" "atomic operat Page 5	of 6
•	Publisher: ACM Press Full text available: pdf(413.70 KB) Additional Information: full citation, index terms	
77	Replication in the harp file system Barbara Liskov, Sanjay Ghemawat, Robert Gruber, Paul Johnson, Liuba Shrira September 1991 ACM SIGOPS Operating Systems Review, Proceedings of the thirteenth ACM symposium on Operating systems principles SOSP '91, Volume 25 Issue 5 Publisher: ACM Press	
	Full text available: pdf(1.60 MB) Additional Information: full citation, abstract, references, citings, index terms	
	This paper describes the design and implementation of the Harp file system. Harp is a replicated Unix file system accessible via the VFS interface. It provides highly available and reliable storage for files and guarantees that file operations are executed atomically in spite of concurrency and failures. It uses a novel variation of the primary copy replication technique that provides good performance because it allows us to trade disk accesses for network communication. Harp is intended to be u	
78	XL-AOF: lightweight aspects for space-based computing	
•	eva Kühn, Fabian Schmied November 2005 Proceedings of the 1st workshop on Aspect oriented middleware development AOMD '05 Publisher: ACM Press	
	Full text available: pdf(395.91 KB) Additional Information: full citation, abstract, references, index terms	
	Space-based computing is a powerful model of abstraction for distributed application development. Although such applications solve a high number of cross-cutting concerns, there is no aspect-oriented environment available at the moment which supports the space-based communication paradigm. This paper describes XL-AOF, an extensible lightweight aspect-oriented framework, whose main focus is to allow for easy development of space-based applications.	
	Keywords : .NET, aspect-oriented programming, declarativity, interception, middleware, space-based computing	
79	Real-Time Refinement and Simplification of Adaptive Triangular Meshes Vasily Volkov, Ling Li October 2003 Proceedings of the 14th IEEE Visualization 2003 (VIS'03) VIS '03 Publisher: IEEE Computer Society	

Full text available: pdf(842.85 KB) Additional Information: full citation, abstract, citings

In this paper we present a generic method for incremental mesh adaptation based on hierarchy of semi-regular meshes. Our method supports any refinement rule mapping vertices onto vertices such as 1-to-4 split or \sqrt3-subdivision. Resulting adaptive mesh has subdivision connectivity and hence good aspect ratio of triangles. Hierarchic representation of the mesh allows incremental local refinement and simplification operations exploiting frame-to-frame coherence. We also present an out-of-core s ...

Keywords: adaptive meshes, refinement and simplification, subdivision, multiresoluton, level of detail, frame-to-frame coherence, out-of-core visualization

80 Fastpath Optimizations for Cluster Recovery in Shared-Disk Systems

Randal Burns

November 2004 Proceedings of the 2004 ACM/IEEE conference on Supercomputing SC '04

Publisher: IEEE Computer Society

Full text available: pdf(176.70 KB) Additional Information: full citation, abstract

We describe the design and implementation of a clustering service for a high-performance, shared-disk file system. The service provides failure detection and recovery, reliableend-to-end messaging, and a centralized and recoverable management interface. We implement novel optimizations in the voting protocol that resolves cluster membership. Optimizations allow clusters to form as quickly as possible without introducing livelock or requiring timeout parameters to be tuned carefully. Our treatmen ...

Results 61 - 80 of 81

Result page: previous 1 2 3 4 5

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

<u>Terms of Usage Privacy Policy Code of Ethics Contact Us</u>

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



• The ACM Digital Library • O The Guide Search:

"data replication" "crash recovery" "atomic operation" "in-men





Feedback Report a problem Satisfaction survey

Terms used data replication crash recovery atomic operation in memory index

Found **81** of **201,890**

Sort results by Display

results

• relevance expanded form

Save results to a Binder Search Tips Open results in a new

Try an Advanced Search Try this search in **The ACM Guide**

window

Results 81 - 81 of 81

Result page: previous 1 2 3 4 5

Relevance scale

Paradigms for process interaction in distributed programs

Gregory R. Andrews

March 1991 ACM Computing Surveys (CSUR), Volume 23 Issue 1

Publisher: ACM Press

Full text available: pdf(3.77 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Distributed computations are concurrent programs in which processes communicate by message passing. Such programs typically execute on network architectures such as networks of workstations or distributed memory parallel machines (i.e., multicomputers such as hypercubes). Several paradigms—examples or models—for process interaction in distributed computations are described. These include networks of filters, clients, and servers, heartbeat algorithms, probe/echo algorithms, broa ...

Keywords: clients and servers, distributed and parallel algorithms, distributed programming, distributed programming methods, heartbeat algorithms, networks of filters, patterns for interprocess communication, probe/echo algorithms, replicated servers, token-passing algorithms

Results 81 - 81 of 81

Result page: previous 1 2 3 4 5

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player